

Weighted Proportional Allocation

Thành Nguyen*
Northwestern University
Evanston, IL, USA
thanh@eecs.northwestern.edu

Milan Vojnović
Microsoft Research
Cambridge, UK
milanv@microsoft.com

ABSTRACT

We consider a weighted proportional allocation of resources that allows providers to discriminate usage of resources by users. This framework is a generalization of well-known proportional allocation by accommodating allocation of resources proportional to weighted bids or proportional to submitted bids but with weighted payments.

We study a competition game where *everyone is selfish*: providers choose user discrimination weights aiming at maximizing their individual *revenues* while users choose their bids aiming at maximizing their individual payoffs. We analyze revenue and social welfare of this game. We find that the revenue is lower bounded by $k/(k+1)$ times the revenue under standard price discrimination scheme, where a set of k users is excluded. For users with linear utility functions, we find that the social welfare is at least $1/(1+2/\sqrt{3})$ of the maximum social welfare (approx. 46%) and that this bound is tight. We extend this efficiency result to a broad class of utility functions and multiple competing providers. We also describe an algorithm for adjusting discrimination weights by providers without a prior knowledge of user utility functions and establish convergence to equilibrium points of the competition game.

Our results show that, in many cases, weighted proportional sharing achieves competitive revenue and social welfare, despite the fact that everyone is selfish.

Categories and Subject Descriptors

K.6.0 [Management of Computing and Information Systems]: General—*Economics*; C.2.3 [Network Architecture Design]: Network Operations

General Terms

Algorithms, Economics, Theory

*Work performed in part while an intern with Microsoft Research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGMETRICS'11, June 7–11, 2011, San Jose, California, USA.
Copyright 2011 ACM 978-1-4503-0262-3/11/06 ...\$10.00.

1. INTRODUCTION

Auction-based Resource Allocation. Use of *pay-per-use pricing* to offer communication and computer services has much proliferated over recent years. For example, cloud computing services are offered through using either fixed pricing or auctions to sell compute instances. The two forms of sales have their own advantages and disadvantages. While fixed pricing schemes are simple to implement, in many scenarios they are neither robust nor flexible enough. For instance, when the user demand is inelastic, a small change in price can translate to a dramatic change in the demand, which can cause congestion and system failure. Furthermore, in order to update prices, providers usually need to gather enough data from sales and this inflexibility can cause system inefficiency and result in low revenue. In recent years, using auction-based schemes for allocating and selling resources in computing systems has become more popular. Sales using auctions are known to be more flexible and can extract information from users. The fact that users can adjust their bids and providers can change the auction parameters in a dynamic fashion yields tremendous improvements of the system efficiency or the revenue extracted by providers. Moreover, the disadvantages of the auction-based approaches such as the difficulty for users to find optimal strategies and for providers to change their platforms have been greatly improved with availability of software that helps users to optimize their bidding strategies and several online services offering auction platforms.

Examples of using auctions for resource allocation include selling of Amazon EC2 spot instances [1], selling of *sponsored search ad slots* to advertisers by major providers of online services, and are of wide interest across engineering systems, e.g. the *electricity markets* [23] that have been of raising interest due to variable supply from renewable energy sources. Furthermore, numerous auction-based allocations of system resources have been proposed such as allocation of disk I/O in storage systems [6], allocation of computational resources [3], and it was even showed that sharing of the Internet bandwidth by TCP connections may be seen as an auction [11, 9].

User Differentiation. In practice providers often apply different prices to different users for selling identical goods or services, which is commonly referred to as *price discrimination*. It is not surprising that similar discrimination schemes are also extensively used in auction-based allocations. There are two main reasons for using such a discriminative framework. First, different users may require different subsets

of resources owned by the provider, e.g. bandwidth over different paths in a communication network, and thus, the provider may improve the system efficiency by differentiating users. Second, different users may have different valuations for the amount of resource received, e.g. different valuation for the amount of network bandwidth received; in this case, once having learned this information, the provider may naturally try to take advantage to increase the revenue by differentiating users.

The most well known example of an auction that differentiates users is the generalized second price auction that is in common use by search engines for allocating ads. In this mechanism, providers (search engines) assign different weights to different users (advertisers) and the mechanism is run based on weighted bids.

The Framework. In this paper, we consider a class of auctions that allows for user differentiation. Specifically, we are interested in auctions that are simple in terms of the information provided by users and are easy to describe to users. We consider two natural instances of weighted proportional allocation: (1) *weighted bid auction* where for a user the allocation is proportional to the product of the bid submitted by this user and the user-specific discrimination weight selected by the provider, and the payment is equal to the bid; (2) *weighted payment auction* where the allocation to a user is proportional the bid of this user and the payment is equal to the product of the bid of this user and the user-specific discrimination weight selected by the provider. The weighted bid auction is a novel proposal while the weighted payment auction was previously considered in Ma et al [15].

As standard in the network pricing literature [18], we consider these allocation mechanisms in the *full information* setting. The justification lies in the fact that in practice allocation auctions are run repeatedly and, thus, providers can learn about the behavior and private information of users. As discussed in the beginning of this section, even in this setting there are advantages of using proportional sharing-like auctions over fixed price schemes. Both auctions that we consider are akin and natural generalizations of the well-known proportional allocation (e.g. [11, 9, 7], see related work discussed later in this section). Therefore, this class of auctions inherits many natural properties of the traditional proportional sharing rule, making it *easy and robust to implement* in practice. Specifically, these mechanisms are simple for bidders as they only need to know the total of others' bids and the allocation is a continuous function of the bids which facilitates robust implementation in a distributed system as will show later in the paper.

Another important reason that motivates us to study the weighted proportional allocation rules is the fact that in settings where provider's goal is to *maximize revenue*, the weighted proportional sharing is preferred over traditional proportional sharing. As it will be shown later, while weighted proportional sharing can provide near-optimal revenue, this is not guaranteed by traditional proportional sharing, which in this regard can be arbitrarily bad.

We study the allocation in *general convex environments* that capture many special cases of resource constraints including those of communication networks, sponsored search, and scheduling of computing resources (see Figure 1 for an illustration).

We describe the weighted proportional schemes consid-

ered in this paper in more detail as follows. We consider a single provider and a set of $n \geq 1$ users and denote with $U = \{1, 2, \dots, n\}$ the set of users. The vectors of allocations and payments are denoted by $\vec{x} = (x_1, x_2, \dots, x_n)$ and $\vec{q} = (q_1, q_2, \dots, q_n)$, respectively. The resource owned by the provider is an arbitrary infinitely divisible resource with the constraints specified by the convex set $\mathcal{P} \in \mathbb{R}_+^n$. An allocation vector \vec{x} is said to be *feasible* if and only if $\vec{x} \in \mathcal{P}$. The provider assigns *discrimination weight* $C_i \geq 0$ to each user i . We denote with $w_i \geq 0$ the bid submitted by user i .

The weighted bid auction is specified by the following allocation and payment rules:

WEIGHTED BID AUCTION

For user i with bid w_i :

$$\text{Allocation} \quad x_i = C_i \frac{w_i}{\sum_{j \in U} w_j}$$

$$\text{Payment} \quad q_i = w_i$$

where the discrimination weights \vec{C} are chosen such that \vec{x} is feasible. We may interpret the discrimination weight C_i as the maximum allocation that is assigned by the provider to user i and x_i as the actual allocation that is the fraction $w_i / \sum_{j \in U} w_j$ of the user-specific maximum allocation C_i .

In turn, the weighted payment auction is specified by the allocation and payment rules defined as follows:

WEIGHTED PAYMENT AUCTION

For user i with bid w_i :

$$\text{Allocation} \quad x_i = C \frac{w_i}{\sum_{j \in U} w_j}$$

$$\text{Payment} \quad q_i = C_i w_i$$

where C is the maximum allocation that is assigned to every user and the discrimination weight C_i determines the payment by user i .

Compared with the traditional proportional allocation, the weighted bid auction is more suitable for general convex resource constraints. This is not the case for the weighted payment auction: while the relative allocation across users can be arbitrary by appropriate choice of the user bids, the implicit assumption of the allocation rule is that $\sum_i x_i = C$, i.e. the provider is required to a priori commit to allocate the total amount of resource C . While this is not restrictive for allocating an infinitely divisible resource of capacity C , this allocation rule cannot accommodate more general polyhedral constraints. Thus, in this paper we will mainly focus on the weighted bid auction but also consider some properties of the weighted payment auction, as it is an alternative auction that allows for user discrimination, albeit for special type of resource constraints.

Questions Studied in this Paper. We consider a competitive setting with multiple providers and users where *everyone is selfish*: each provider aims at maximizing own revenue and each user aims at maximizing own payoff. Ideally, an allocation mechanism would guarantee high *revenue* to the provider and high *efficiency* where by efficiency we mean social welfare (i.e. the total utility across all users) compared with the best possible social welfare. In the competitive setting where everyone is selfish, it is rather unclear whether the two goals could be achieved simultaneously. Intuitively, one would expect that selfishness of providers and




Figure 1: Examples of polyhedron constraints: (left) single link, (middle) a network of links, (right) assignment costs of a workflow to machines.

users may well result in either low revenue or efficiency. For example, providers that aim at maximizing their revenue may well have an incentive to misreport availability of their resources¹. The main questions that we consider in this paper are:

- Q1: *How much of the revenue can be guaranteed using the weighted proportional sharing?*
- Q2: *What is the efficiency loss in competitive systems where everyone is selfish?*

Summary of Results. We summarize our results in the following points:

- *Revenue:* We show that the revenue of weighted bid auction is at least $k/(k+1)$ times the revenue under standard price discrimination scheme with a set of k users excluded, which we describe in more detail in Section 3. The comparison of the revenue of a mechanism with the maximum revenue obtained under another mechanism where some users are excluded is standard in the mechanism design literature (e.g. [8]) and in this context, our revenue comparison result is novel and may be of general interest. The result enables us to understand conditions under which the revenue of weighted bid auction is competitive to that under the benchmark pricing scheme, e.g. the case of many users.

- *Efficiency:* We establish that for linear user utility functions, the social welfare under weighted bid auction is at least $1/(1+2/\sqrt{3}) \approx 0.46$ -factor of the maximum social welfare and this bound is tight. We also show a tight efficiency bound of $1/2$ for the weighed payment auctions. We extend our efficiency result of the weighed bid auction to a broad class of utility functions, we call δ -utility functions, where $\delta \geq 0$. We show that many utility functions found in literature are δ -utility functions, in many cases with $\delta \leq 2$, and show that this class of utility functions is closed to addition and multiplication with a positive constant. For the class of δ -utility functions, we show that the social welfare is at least $1/(1+2/\sqrt{3}+\delta)$ -factor of the maximum social welfare and establish that this guarantee holds also for the case of multiple competing providers. A similar extension for the weighted payment auction is also shown.

¹A famous example of such a market manipulation is *California electricity crises* where in 2000 and 2001 there was a shortage of electricity because energy traders were gratuitously taking their plants offline at peak demand in order to sell at higher prices [24].

- *Convergence and distributed algorithms:* We describe how the provider can adjust discrimination weights in an online fashion so that user allocations and payments converge to equilibrium points of the underlying competition game. This shows that the information about user utility functions needed by the provider can be estimated from the bids submitted by users in an online fashion. We describe a distributed iterative scheme and prove convergence to Nash equilibrium points for the case of linear user utility functions.

Related Work. The allocation of resources in proportion to user-specific weights has a long and rich history in the context of computer systems and services. For example, it underlies the objective of generalized processor sharing [19, 2], sharing of bandwidth in communication networks [11], and has been considered for allocation of various types of resources, including storage [6] and compute instances [3]. The weighted bid auction considered in this paper could be seen as a generalization of traditional proportional allocation to allow for user differentiation and general convex resource constraints.

Previous work primarily focused on analyzing social efficiency of proportional allocation in competitive environments where only users are assumed to be selfish. Kelly [11] showed that under price taking users, where each user submits a scalar bid for a set of resources of interest (e.g. allocation of bandwidth at each link along the path of a network connection), the proportional allocation guarantees 100% efficiency. In subsequent work, Johari and Tsitsiklis [9] showed that the social efficiency is at least 75% under so called price anticipating users and assumption that each user submits an individual bid for each individual resource of interest (e.g. an individual bid submitted for each link along the path of a network connection). The latter result was extended by Nguyen and Tardos [16] to more general polyhedral constraints. Furthermore, Yang and Hajek [7] showed that the worst-case efficiency is 0% if users are restricted to submitting a scalar bid for the set of resources of interest. Our work is different from all this work in that we consider a competitive environment where everyone is selfish.

The weighted payment auction was considered by Ma et al [15] where the focus was on the efficiency with respect to the social welfare. Our work provides new results for this type of auctions with regard to the revenue and efficiency in competitive environments where everyone is selfish.

The setting of multiple providers considered in this paper bears quite some similarity with that found in the context of

ISP multihoming, e.g. [5, 20], multi-path congestion control, e.g. [4], and may also inform about the competition in time-varying markets using the approach in [13].

Outline of the Paper. Section 2 introduces the resource competition game. Section 3 presents our main revenue comparison result (Theorem 3.1). In Section 4, we present our results on the efficiency guarantees for the case of a single provider and users with linear utility functions (Theorem 4.1). Section 5 extends the efficiency result to more general class of user utility functions and more general setting of multiple providers (Theorem 5.1). We discuss convergence to Nash equilibrium points and distributed schemes in Section 6. In Section 7, we conclude.

2. THE RESOURCE ALLOCATION GAME WHERE EVERYONE IS SELFISH

We consider a system of $n \geq 1$ users competing for resources of a single provider; we introduce the setting with multiple providers later in Section 5.1. Recall that $\vec{x} = (x_1, x_2, \dots, x_n)$ and $\vec{q} = (q_1, q_2, \dots, q_n)$ denote the vectors of allocations and payments by users, respectively. The allocation vector \vec{x} is feasible if and only if $\vec{x} \in \mathcal{P}$ where \mathcal{P} is a convex set of the form $\mathcal{P} = \{\vec{x} \in \mathbb{R}_+^n : A\vec{x} \leq \vec{b}\}$ for some matrix A and vector \vec{b} with non-negative elements. (Note that A and \vec{b} can have arbitrarily many rows, so we will often refer to “convex” constraints instead to polyhedrons.)

Suppose that $U_i(x_i)$ is the utility of allocation x_i to user i . Throughout this paper we assume that for every i , $U_i(x)$ is a non-negative, non-decreasing and continuously differentiable concave function. The payoffs for the provider and users are defined as follows. The payoff of the provider is equal to the revenue, i.e. the total payments received from all users, $R = \sum_i q_i$. The payoff for user i is equal to the utility minus the payment, i.e. $U_i(x_i) - q_i$.

The competition game that we study can be seen as the following *two-stage Stackelberg game*: in the first stage, the provider announces the discrimination weights \vec{C} and then, in the second stage, users adjust their bids in a selfish way aiming at maximizing their individual payoffs. In the first stage, the provider anticipates how users would react to given discrimination weights \vec{C} and sets these weights in a selfish way aiming at maximizing the revenue. In a dynamic setting, the two stages of this game would alternate over time (we discuss this in Section 6).

In the reminder of this section, we characterize the Nash equilibria for weighted-bid and weighted-payment auction.

Equilibrium of Weighted Bid Auction. We show a relation between the revenue and the allocation of an outcome. Given discrimination weight C_i and sum of the bids $\sum_j w_j$, user i selects bid w_i that maximizes his surplus, i.e. solves

$$\text{USER: } \max_{w_i \geq 0} U_i \left(\frac{w_i}{\sum_{j \neq i} w_j + w_i} C_i \right) - w_i. \quad (1)$$

Under the assumed user behavior, one can analyze the Nash equilibrium of the game. It turns out that Nash equilibrium exists and is unique, and at Nash equilibrium, the relation between the revenue and allocation is captured by an implicit function, which we show in the following lemma.

LEMMA 2.1. *Given a vector of discrimination weights \vec{C} , there is a unique allocation \vec{x} corresponding to unique*

Nash equilibrium. Conversely, given an equilibrium allocation vector \vec{x} , there is a unique vector of discrimination weights \vec{C} . Furthermore, the corresponding revenue $R(\vec{x})$ is a function of the allocation vector \vec{x} given by

$$\sum_i \frac{U'_i(x_i)x_i}{U'_i(x_i)x_i + R(\vec{x})} = 1, \quad (2)$$

where $U'_i(x_i)$ is the derivative of U_i at x_i .

Given this result, we obtain the following optimization problem for the provider.

$$\text{PROVIDER: } \text{maximize } R(\vec{x}) \text{ over } \vec{x} \in \mathcal{P}. \quad (3)$$

In the rest of this section we prove Lemma 2.1.

PROOF OF LEMMA 2.1. We have

$$x_i = C_i \frac{w_i}{\sum_j w_j} \quad (4)$$

and USER problem can be written as:

$$\text{maximize } U_i \left(\frac{w_i}{\sum_{j \neq i} w_j + w_i} C_i \right) - w_i \text{ over } w_i \geq 0. \quad (5)$$

Note that the objective function in (5) is concave in w_i , hence, at an optimum solution either $w_i = 0$ or the derivative of the objective function is zero. Setting the derivative to zero is equivalent to:

$$U'_i(x_i) \cdot C_i \frac{\sum_{j \neq i} w_j}{(\sum_j w_j)^2} = 1, \text{ for } x_i > 0.$$

It follows

$$U'_i(x_i) = \frac{(\sum_j w_j)^2}{C_i \sum_{j \neq i} w_j} = \frac{R^2}{C_i(R - w_i)} \quad (6)$$

where recall that the revenue is equal to the sum of payments made by users, i.e. $R = \sum_j w_j$. Combining with $w_i = x_i R / C_i$, which follows from (4), we have

$$U'_i(x_i) = \frac{R}{C_i - x_i} \Leftrightarrow C_i U'_i(x_i) \left(1 - \frac{x_i}{C_i} \right) = R. \quad (7)$$

Now, $\sum \frac{x_i}{C_i} = 1$, thus, condition (7) is exactly the optimality condition for the following problem

$$\begin{aligned} & \text{maximize } \sum_i \int_0^{x_i} C_i U'_i(t_i) \left(1 - \frac{t_i}{C_i} \right) dt_i \\ & \text{over } \vec{x} \in \mathbb{R}_+^n \\ & \text{subject to } \sum_i \frac{x_i}{C_i} = 1. \end{aligned}$$

Since $\int_0^{x_i} C_i U'_i(t_i) \left(1 - \frac{t_i}{C_i} \right) dt_i$ is a strictly concave function with respect to x_i , there exists a unique Nash equilibrium.

It remains to show that for an equilibrium allocation \vec{x} , the revenue R is given by

$$\sum_i \frac{U'_i(x_i)x_i}{U'_i(x_i)x_i + R} = 1. \quad (8)$$

From (7), we have

$$U'_i(x_i) = \frac{R}{C_i - x_i} \Rightarrow \frac{C_i}{x_i} - 1 = \frac{R}{U'_i(x_i)x_i}$$

$$\Rightarrow \frac{x_i}{C_i} = \frac{U'_i(x_i)x_i}{U'_i(x_i)x_i + R}.$$

Combining with $\sum_i x_i/C_i = 1$, which follows from (4), we obtain (8). Note that all the formulas above are applied for the case $x_i > 0$ only; nevertheless, if $x_i = 0$, we have $U'_i(x_i)x_i = 0$, and therefore, the equation (8) holds for any optimum allocation vector \vec{x} .

Finally, we note that in equilibrium, the vector of discrimination weights \vec{C} and the vector of bids \vec{w} are functions of the equilibrium allocation \vec{x} given as follows: for every i ,

$$C_i = x_i + \frac{R(\vec{x})}{U'_i(x_i)} \text{ and } w_i = \frac{R(\vec{x})}{U'_i(x_i)x_i + R(\vec{x})} U'_i(x_i)x_i.$$

□

Equilibrium of Weighted Payment Auction. The analysis follows similar steps as for weighted bid auction. In this case, the revenue in Nash equilibrium can be represented as an explicit function of the allocation vector \vec{x} . Given discrimination weight C_i , user i solves the following surplus maximization problem:

$$\text{USER: } \max_{w_i \geq 0} U_i \left(\frac{w_i}{\sum_{j \neq i} w_j + w_i} C \right) - C_i w_i. \quad (9)$$

LEMMA 2.2. *Given a vector of discrimination weights \vec{C} , there is a unique allocation \vec{x} corresponding to the unique Nash equilibrium. Conversely, given an allocation \vec{x} , such that $\sum_i x_i = C$, there is a vector \vec{C} of discrimination weights such that \vec{x} is an outcome. Furthermore, the corresponding revenue $R(\vec{x})$ is given by*

$$R(\vec{x}) = \sum_i U'_i(x_i) \frac{x_i}{C} \left(1 - \frac{x_i}{C} \right).$$

The proof of this lemma is provided in [17].

3. REVENUE

Revenue of Proportional Sharing. We provide an example showing that traditional proportional sharing can perform poorly with respect to the revenue. The example is for the parking-lot network scenario that has been extensively used in the context of networking; see Figure 2 for an illustration. The resource consists of a series of $n \geq 1$ links, each of capacity $C > 0$ (without loss of generality we assume $C = 1$). The example consists of $n + 1$ users; user 0 is a *multi-hop user* that requires a connection through all links $1, 2, \dots, n$ while user i is a *single-hop user* that requires a connection through link i . User utility functions are assumed to be α -fair, i.e. for $\alpha > 0$ and $\alpha \neq 1$, we have $U_i(x) = \frac{\gamma_i^\alpha}{1-\alpha} x^{1-\alpha}$, and $U_i(x) = \gamma_i \log(x)$, for $\alpha = 1$, where $\gamma_i > 0$ (in our example, we consider symmetric case where $\gamma_i = 1$, for every user i). The resource constraints in this example are $x_i + x_0 \leq 1$, for every link i .

Under proportional sharing mechanism, user 0 submits an individual bid for each link while user $i > 0$ only bids for link i . Using the known conditions for Nash equilibrium of the underlying game (e.g. [9, 16]), one can show that there is a unique equilibrium with the allocation vector, solution




Figure 2: Parking-lot example.

of the following problem:

$$\begin{aligned} & \text{maximize } \sum_{i=0}^n \int_0^{x_i} U'_i(y)(1-y)dy \\ & \text{over } \vec{x} \in \mathbb{R}_+^{n+1} \\ & \text{subject to } x_i + x_0 \leq 1, \quad i = 1, 2, \dots, n. \end{aligned}$$

It is known that in Nash equilibrium, the sum of bids submitted to link i is equal to $U'_i(x_i)(1-x_i)$ [16]. Therefore, the revenue of proportional sharing in Nash equilibrium is $\sum_{i=1}^n U'_i(x_i)(1-x_i)$. By straightforward calculations, one can show that in Nash equilibrium, the allocation is x_0 to user 0 and $1-x_0$ to each user $i > 1$, where $x_0 = \frac{1}{1+n^{1/(\alpha+1)}}$, and the revenue is $\frac{nx_0}{(1-x_0)^\alpha}$. For large network size n , the revenue in Nash equilibrium is $O(n^{\frac{\alpha}{\alpha+1}})$, therefore, $o(n)$, for every fixed $\alpha > 0$. The intuition behind why the revenue is low is: for large network size n , user 0 competes against many users and, therefore, the payment that this user can afford at each link is small. Consequently, the competition at each link is low which results in low revenue.

On the other hand, using fixed pricing, the provider can charge the price per unit resource of 1 on each link and receive the revenue of n .

The question that we investigate in this section is how large revenue the weighted bid auction can achieve, in comparison with the best fixed pricing scheme (with price discrimination). We will answer this question in the remainder of this section.

Revenue of Weighted Bid Auction. We will compare the revenue of weighted bid auction with that of a benchmark that consists of standard fixed price scheme [22]. In this pricing scheme, the provider charges user-specific prices per unit of resource for different users. Suppose that provider charges user i , the price per unit resource p_i . Then, user i surplus maximization problem is $\max U_i(x_i) - p_i x_i$ over $x_i \geq 0$. The solution of this problem is given by $U'_i(x_i) = p_i$. Therefore, the revenue of the provider is $R = \sum_i p_i x_i = \sum_i U'_i(x_i)x_i$. Hence, the optimal revenue is

$$R^* = \max \left\{ \sum_i U'_i(x_i)x_i : \vec{x} \in \mathcal{P} \right\}.$$

Comparing with such a benchmark would be too ambitious because with auction-based allocation, the provider cannot announce fixed prices, but instead prices are induced from user demand. Thus, instead, we will compare with the revenue R^* where *some users are excluded*. That is, we will compare the revenue achieved by an auction for n users with the revenue achieved by the fixed pricing scheme for $n-k$ users, for $0 < k < n$. Note that using as a benchmark a scheme with some users excluded is standard in the theory of auctions [8].

In the parking-lot example introduced above, if we con-

sider n large enough, then the optimal revenue is of the order n , which can be achieved if the provider charges every single-hop user the price per unit of resource of 1 and charges the multi-hop user the price per unit of resource of n . Now, if we exclude an arbitrary set of $k < n$ users, then the optimal revenue is $n - k$. Therefore, if k is much smaller compared with n , then one can think of this as a *large market* regime where the effect of removing a few users from the market is negligible.

Having discussed the intuition, we can now state our main result on the revenue guarantee of weighted bid auctions. Let R_{n-k}^* be the optimal revenue under our benchmark, i.e.

$$R_{n-k}^* = \min_{S \subset \{1, \dots, n\}: |S|=n-k} \max_{\vec{x} \in \mathcal{P}} \sum_{i \in S} U_i'(x_i) x_i.$$

The revenue guarantee of weighted bid auctions is stated as follows.

THEOREM 3.1. *Suppose that for each user i , $U_i'(x)$ is a concave function. Let R be the optimum revenue of the weighted bid allocation mechanism, then*

$$\text{for every } 1 \leq k < n: R \geq \frac{k}{k+1} R_{n-k}^*.$$

The proof of the theorem is based on an induction argument over k and is provided in [17]. It is noteworthy that the proof admits weak assumptions about the structure of the underlying auction. For example, using an analogue argument, we can also prove similar revenue guarantee for weighted payment auctions, which is provided in [17].

Finally, note that the revenue guarantee of the theorem above is rather strong. In particular, by taking as a benchmark the system with just one user excluded, we obtain that the revenue under weighted bid auction is at least 1/2 of the revenue under the price discrimination scheme with one user excluded (whose exclusion reduces the revenue the most). Informally, the result tells us that for systems with many users with comparable utility functions, the revenue under weighted bid auction is nearly the same as under standard price discrimination. As discussed above, such a guarantee cannot be provided by proportional sharing.

4. EFFICIENCY FOR LINEAR USER UTILITY FUNCTIONS

In this section, we analyze the efficiency for the case of single provider and linear user utility functions. The analysis in this section provides us with basic techniques that are applied to the more general setting in Section 5.

4.1 Efficiency of Weighted Bid Auction

We show the following theorem.

THEOREM 4.1. *Assume that the provider maximizes the revenue and for each user i the utility function is linear, $U_i(x) = v_i x$, for some $v_i > 0$. Then, the worst-case efficiency is $1/(1 + 2/\sqrt{3})$ (approx. 46%). Furthermore, this bound is tight.*

Remark Before proving the theorem, we note that the worst-case efficiency can be achieved asymptotically as the number of users n tends to infinity. An example is for the resource constraint $\sum_i x_i \leq 1$ and the user valuations such that there is a unique user with the largest marginal utility, say this is

user 1, and all other users with identical marginal utilities equal to $(2 - \sqrt{3})^2 v_1 \approx 0.0718 v_1$. In Nash equilibrium, user 1 obtains 42.26% of the resource and the rest is equally shared by other users. One can show that with higher user competitiveness, the efficiency increases. Specifically, if there are at least k users with the largest marginal utility, then the efficiency is at least $1 - \frac{1}{2k} + o(1/k)$. The proof of this result is in [17].

We first need the following lemma about the *quasi-concavity* of the objective function optimized by the provider. Recall that $R(\vec{x})$ is the function given by (2). Let R^* be the optimum revenue, i.e. $R^* = \max\{R(\vec{x}) : \vec{x} \in \mathcal{P}\}$. We note the following fact.

LEMMA 4.1. *The set $\mathcal{L}_\mu := \{\vec{x} \in \mathbb{R}_+^n : R(\vec{x}) \geq \mu\}$ is convex, for every $\mu \in [0, R^*]$.*

The proof of this lemma is straightforward and thus omitted (see [17]). In the remainder of this section, we prove Theorem 4.1.

PROOF OF THEOREM 4.1. An example showing that the bound is tight was already given in the remark above; it remains to prove that the efficiency is at least $1/(1 + 2/\sqrt{3})$.

Recall that R^* is the optimal revenue, thus for every $\vec{x} \in \mathcal{P}$, $R(\vec{x}) \leq R^*$. Consider the two convex sets \mathcal{L}_{R^*} and \mathcal{P} . These two sets intersect at \vec{x} where $R(\vec{x}) = R^*$ and do not have common interior points. Let H be the hyperplane that weakly separates these two sets, defined by $\gamma_i \geq 0$, for every i and

$$\sum_i \gamma_i x_i = 1. \quad (10)$$

Consider the game where the provider has the feasible set $\mathcal{Q} = \{\vec{x} \in \mathbb{R}_+^n : \sum_i \gamma_i x_i \leq 1\}$, then the allocation that maximizes the revenue over \mathcal{Q} is the same as in the original game. Since $\mathcal{P} \subset \mathcal{Q}$, the optimal social welfare of the new game is at least the social welfare of the original game. Therefore, it is enough to prove a lower bound on the efficiency for the class of games where the provider has the feasible set \mathcal{Q} . See Figure 3 for an illustration.




Figure 3: Reduction to a simpler constraint.

The observation above allows us to consider a simpler optimization problem. In particular, the optimal social welfare in the new game is $\max_i v_i/\gamma_i$; the condition for Nash equilibrium, as argued above, is the condition for \vec{x} to maximize $R(\vec{x})$ over $\vec{x} \in \mathbb{R}_+^n$ such that $\sum_i \gamma_i x_i = 1$, which we derive in the following. Taking partial derivative with respect to x_j to both sides in (2), with $U_i(x_i) = v_i x_i$, we have

$$\begin{aligned} \frac{\partial}{\partial x_j} \sum_i \frac{v_i x_i}{v_i x_i + R} &= 0 \Leftrightarrow \\ \Leftrightarrow \frac{\partial}{\partial x_j} \frac{v_j x_j}{(v_j x_j + R)} - \sum_{i \neq j} \frac{\partial R}{\partial x_j} \frac{v_j x_i}{(v_i x_i + R)^2} &= 0. \end{aligned}$$

Now, note that

$$\frac{\partial}{\partial x_j} \frac{v_j x_j}{v_j x_j + R} = \frac{R v_j}{(v_j x_j + R)^2} - \frac{\partial R}{\partial x_j} \frac{v_j x_j}{(v_j x_j + R)^2}.$$

Thus, we have

$$\frac{R v_j}{(v_j x_j + R)^2} = \frac{\partial R}{\partial x_j} \cdot \sum_i \frac{v_i x_i}{(v_i x_i + R)^2}.$$

Since $R(\vec{x})$ achieves the optimum value R^* over the set $\{\vec{x} \in \mathbb{R}_+^n : \sum_i \gamma_i x_i \leq 1\}$, we have either $x_j = 0$ or $\frac{\partial}{\partial x_j} R = \lambda \gamma_j$ where $\lambda \geq 0$ is the Lagrange multiplier associated to the constraint $\sum_i \gamma_i x_i \leq 1$. It follows that for $p > 0$,

$$\begin{aligned} &\text{either } x_i = 0 \\ &\text{or } \frac{v_i/\gamma_i}{(v_i x_i + R^*)^2} = \frac{\lambda}{R^*} \sum_i \frac{v_i x_i}{(v_i x_i + R^*)^2} = p. \end{aligned} \quad (11)$$

From this, we obtain that at in Nash equilibrium, if $x_i > 0$, then $\frac{v_i/\gamma_i}{(v_i x_i + R^*)^2}$ is equal to p , for every i . Therefore, if v_i/γ_i is large, then the denominator $(v_i x_i + R^*)^2$ needs to be large as well. At the same time, the optimal solution of social welfare distributes all the resource to users with the highest value of v_i/γ_i . This is the intuition for the fact that the efficiency is bounded by a constant.

First, we will use the change of variables to make equations easier to follow: $z_i = \gamma_i x_i$ and $a_i = v_i/\gamma_i$. One way to think about these new variables is to think of another game where the resource constraint is $\sum_i z_i = 1$ and user i 's utility is $a_i z_i$. Without loss of generality, we assume that $a_1 = \max_i a_i$. The optimal social welfare is

$$W_{\text{OPT}} = \max_{\vec{z}: \sum_i \gamma_i x_i = 1} \sum_i v_i x_i = \max_{\vec{z}: \sum_i z_i = 1} \sum_i a_i z_i = a_1.$$

We also introduce variables y_i defined by

$$y_i = \frac{v_i x_i}{v_i x_i + R^*} = \frac{a_i z_i}{a_i z_i + R^*}.$$

Because of (2), we have $\sum_i y_i = 1$. The goal of introducing these variables is to bound the optimal social welfare and the social welfare in Nash equilibrium as functions of y_i . From $y_i = a_i z_i / (a_i z_i + R^*)$, we have

$$a_i z_i = R^* \frac{y_i}{1 - y_i} \quad \text{and} \quad z_i = R^* \frac{y_i}{a_i(1 - y_i)}.$$

Next, we are going to bound the social welfare of Nash equilibrium and the optimal solution.

The social welfare in Nash equilibrium, which we denote as W_{NASH} , can be bounded as follows

$$\begin{aligned} W_{\text{NASH}} &= \sum_i a_i z_i = R^* \sum_i \frac{y_i}{1 - y_i} \\ &\geq R^* \left(\frac{y_1}{1 - y_1} + \sum_{i \geq 2} y_i \right) = R^* \left(\frac{y_1}{1 - y_1} + 1 - y_1 \right) \\ &\geq R^* \frac{y_1^2 - y_1 + 1}{1 - y_1}. \end{aligned} \quad (12)$$

The optimal social welfare, as argued above, is

$$W_{\text{OPT}} = \max_i a_i = a_1.$$

In order to bound a_1 with a function of y_i , we multiply a_1 with $\sum_i z_i$, which is 1, and use the relation between z_i and

y_i to have W_{OPT} as a function of y_i . Specifically,

$$W_{\text{OPT}} = a_1 = a_1 \left(\sum_i z_i \right) = a_1 R^* \sum_i \frac{y_i}{a_i(1 - y_i)}. \quad (13)$$

Now, we use the condition for Nash equilibrium. (Note that this is the only place in the proof that uses (11).) First we rewrite the condition for the variables z_i and a_i . Replacing $a_i = v_i/\gamma_i$ and $v_i x_i = a_i z_i = R^* \frac{y_i}{1 - y_i}$ in the condition for Nash equilibrium (11), we derive

$$\text{either } y_i = 0 \text{ or } \frac{a_i(1 - y_i)^2}{R^{*2}} = p > 0.$$

From this condition, we have $a_i(1 - y_i)^2 = a_1(1 - y_1)^2$ whenever $y_1, y_i > 0$, hence $a_i(1 - y_i) = \frac{a_i(1 - y_1)^2}{1 - y_i}$. Replacing this equality in the optimal social welfare (13), we have

$$\begin{aligned} W_{\text{OPT}} &= a_1 R^* \sum_i \frac{y_i}{a_i(1 - y_i)} = \frac{R^*}{(1 - y_1)^2} \sum_i y_i(1 - y_i) \\ &\leq \frac{R^*}{(1 - y_1)^2} \left(y_1(1 - y_1) + \sum_{i \geq 2} y_i \right). \end{aligned}$$

Using this and replacing $\sum_{i \geq 2} y_i = 1 - y_1$, we obtain

$$W_{\text{OPT}} \leq \frac{R^*}{(1 - y_1)^2} (y_1(1 - y_1) + 1 - y_1) = R^* \frac{1 - y_1^2}{(1 - y_1)^2}. \quad (14)$$

From (12) and (14), we have the following lower bound for the efficiency

$$\frac{W_{\text{NASH}}}{W_{\text{OPT}}} \geq \frac{y_1^2 - y_1 + 1}{y_1 + 1}.$$

By simple calculus, one can show that the right-hand side is at least $1/(1 + 2/\sqrt{3})$, which is what we needed to prove. \square

4.2 Efficiency of Weighted Payment Auction

For weighted payment auction, we have the following result.

THEOREM 4.2. *Assuming that the provider maximizes the revenue, the efficiency is at least 1/2 for weighted payment auction and this bound is tight.*

The proof of this theorem is provided in [17]. We note that the weighted payment auction admits much simpler resource constraint and, therefore, the proof of this theorem is much simpler than that for weighted bid auctions, which we presented above.

5. EXTENSION TO MULTIPLE PROVIDERS AND GENERAL USER UTILITIES

In this section, we will extend the efficiency result of the previous section to the more general setting that consists of multiple competing providers and more general class of utility functions. Perhaps surprisingly, we will show that even in such more complex competitive environments, the efficiency can be bounded by a positive constant that is independent of the number of providers and the number of users. We first define the framework in Section 5.1 and Section 5.2, and then present our main result in Section 5.3.

5.1 Multiple Providers

We consider a system of multiple providers where each provider allocates resources according to the weighted proportional allocation. We assume that each provider k is endowed with resource constraints specified by the convex set \mathcal{P}_k . We assume that each user can receive resources from any provider and is concerned only about the total amount of resource received across all providers. We will use the following notation. We denote with x_i^k the allocation to user i by provider k . Let $x_i = \sum_k x_i^k$ denote the total allocation to user i over all providers. For each user i , the utility of allocation $(x_i^k, k = 1, \dots, m)$ is $U_i(x_i)$. We denote with $x_i^{-k} = x_i - x_i^k$, the total allocation to user i over all providers except provider k . See Fig. 4 for an illustration.




Figure 4: The setting of multiple providers.

Let $\vec{x} = (x_i^k, i = 1, \dots, n, k = 1, \dots, m)$ be an allocation under weighted proportional sharing mechanism. By same arguments as in Section 2, we note that each provider k can find discrimination weights $(C_1^k, C_2^k, \dots, C_n^k)$ such that \vec{x} is the equilibrium of the weighted proportional sharing in the multiple provider setting. We denote with w_i^k the bid of user i to provider k and this is the payment from user i to provider k . The user i goal is to maximize the payoff $U_i(\sum_k x_i^k) - \sum_k w_i^k$, where $x_i^k = C_i^k w_i^k / \sum_i w_i^k$. On the other hand, provider k obtains the revenue R^k , which satisfies the following

$$\sum_{i=1}^n \frac{U_i'(x_i^{-k} + x_i^k) x_i^k}{U_i'(x_i^{-k} + x_i^k) x_i^k + R^k} = 1. \quad (15)$$

In order to gain some intuition, note that if for every user i , $U_i(x)$ is a strictly concave function, then $U_i'(x_i^{-k} + x_i^k)$ decreases with x_i^{-k} . In other words, the larger the amount of resource allocated to user i from providers other than k , the smaller the marginal utility for user i to receive an allocation from provider k . As a result, provider k may extract smaller revenue due to competition with other providers. With this in mind, we now define the equilibrium for the case of multiple providers.

DEFINITION 1. We call \vec{x} an equilibrium allocation if for every provider k , the allocation vector $\vec{x}^k = (x_1^k, \dots, x_n^k)$ maximizes R^k given by (15) over the set \mathcal{P}_k .

We note that in the multiple provider setting, we can think of the game as the provider k 's strategy set is \mathcal{P}_k . The discrimination weights and the revenue then can be calculated according to the allocation vector \vec{x} of all providers. With these discrimination weights, under users' selfish behavior, \vec{x} will be an outcome of the game. From the providers' perspective, an equilibrium allocation is an allocation \vec{x} where

no provider has an incentive to unilaterally change its allocation vector. Note that if there is only one provider, the game boils down to the same two-stage Stackelberg game considered in Section 2.

5.2 A Class of Utility Functions

We introduce the class of δ -utility functions in the following definition.

DEFINITION 2. Let $U(x)$ be a non-negative, increasing, and concave utility function and let $x_0 \geq 0$ be the value maximizing $U'(x)x$. We say $U(x)$ is a δ -utility, if in addition the following two conditions hold:

- (i) $U'(x)x$ is a concave function over $[0, x_0]$, and
- (ii) there exists $\delta \in [0, \infty)$, such that, for every $a \in [0, x_0]$,

$$U(b) - [U'(a)a]^b \leq \delta U(a)$$

where $b \geq 0$ is such that $U'(b) = [U'(a)a]^b = U'(a) + U''(a)a \geq 0$.

The class of δ -utility functions is intimately related with the theory of price discrimination [22], which we discuss in [17]. Furthermore, the class accommodates many utility functions considered in literature. For example, we can show that a linear or a truncated linear function is a 0-utility function, polynomial $(c+x)^\alpha$, for $c \geq 0$ is a $\frac{\alpha}{2}$ -utility function, for any $0 \leq \alpha \leq 1$, and a logarithmic utility function is a 2-utility function; we provide a detailed list and proofs in [17]. We remark that truncated linear utility functions or logarithmic functions were considered representative of real-time traffic sources in communication networks [21], polynomial utility functions were widely used in economic theory [22], α -fair utility functions were widely used in the context of communication networks [12, 10].

Finally, we note the following result whose proof is provided in [17].

LEMMA 5.1. If f and g are δ -utility functions, then so are: $c \cdot f$, for $c > 0$ and $f + g$.

Consequently, every polynomial of the form $\sum_i a_i x^{\alpha_i}$ where $a_i > 0$ and $0 \leq \alpha_i \leq 1$ is a $\frac{\alpha}{2}$ -utility function.

5.3 Efficiency Bound

We now state and then prove our main theorem on the efficiency of weighted bid auctions.

THEOREM 5.1. Assume that for every user i and every $a \geq 0$, $U_i'(x+a)x$ is a continuous and concave function. Then, there exists an equilibrium in the case of multiple providers defined as above. Furthermore, if $U_i(a+x)$ are δ -utility functions, then the efficiency at any equilibrium is at least $1/(1+2/\sqrt{3}+\delta)$.

For the special case of linear utility functions $\delta = 0$ and the above theorem yields the same bound as in Theorem 4.1. The result of Theorem 5.1 is perhaps surprising as it is not a priori clear that in complex competitive environments where both providers and users are selfish in trying to maximize their individual payoffs (objectives which often conflict with each other), the efficiency would be bounded by a positive constant that is independent of the number of providers and the number of users.

Before going into the proof of Theorem 5.1, we outline the main ideas. The key idea of the proof is to bound the social welfare of the system, which is a complicated optimization problem over allocations in the Minkowski sum of the sets \mathcal{P}_k . If the utility functions are linear, then this optimization problem can be *separated* into optimization problems for individual providers, and the optimal value is the sum of these optimal values. Using this idea, we will bound the utility function by an affine function that is a tangent to the concave utility function at particular allocation in Nash equilibrium. This idea is illustrated in Figure 5. It will be shown that because of the property of δ -utility functions, the value a_i in the figure is at most $\delta U_i(x_i)$, which will be the key inequality of the proof.²




Figure 5: The key bounding of $U_i(x)$ with the affine function $V_i(x)$ such that $V_i'(y_i) = U_i'(x_i) + U_i''(x_i) \max_k x_i^k$ and $V_i(y_i) = U_i(y_i)$.

PROOF OF THEOREM 5.1. The proof of the first part of the theorem about existence of Nash equilibrium is based on standard fixed point theorem and is provided in [17].

For the second part, the key idea of the proof is to bound the social welfare by an affine function which allows separating the maximization over $(\bar{x}^1, \dots, \bar{x}^m) \in \sum_k \mathcal{P}_k$ to maximizations over sets \mathcal{P}_k , where $\sum_k \mathcal{P}_k$ is the Minkowski sum defined by $\{\bar{z}^1 + \dots + \bar{z}^m : \bar{z}^k \in \mathcal{P}_k, k = 1, \dots, m\}$. Once the optimization problem is separated, we can use similar bound as in Section 4 (see Lemma 5.2 below) as a subroutine to prove the theorem.

Let us define

$$v_i^k = U_i'(x_i) + U_i''(x_i)x_i^k \text{ and } v_i = \min_k v_i^k.$$

Since for every i , $U_i(x)$ is a concave function, we have

$$v_i = U_i'(x_i) + U_i''(x_i)(\max_k x_i^k) \geq U_i'(x_i) + U_i''(x_i)x_i.$$

Now, let us define $V_i(x) = a_i + v_i x$ where a_i is chosen such that $V_i(x)$ is a tangent to $U_i(x)$. Let y_i be the point at $V_i(x)$ and $U_i(x)$ intersect. We will use $V_i(x)$ as an upper bound for $U_i(x)$, for every $x \geq 0$. Notice that $a_i = U_i(y_i) - (U_i'(x_i) + U_i''(x_i)x_i)y_i$.

By definition of δ -utility functions, $a_i \leq \delta U_i(x_i)$. Therefore,

$$\sum_i a_i \leq \delta \sum_i U_i(x_i). \quad (16)$$

²We note that our proof technique is rather general and that similar result for weighted payment auctions can also be obtained using this framework. However, we omit details as our focus is on weighted bid auctions, which allow for much more general resource constraints.

Since $U_i(x)$ is a non-negative concave function, we have $U_i(x) \leq V_i(x)$. Hence,

$$\begin{aligned} \max_{\bar{z} \in \sum_k \mathcal{P}_k} \sum_i U_i(z_i) &\leq \max_{\bar{z} \in \sum_k \mathcal{P}_k} \sum_i V_i(z_i) = \\ &= \sum_i a_i + \max_{\bar{z} \in \sum_k \mathcal{P}_k} \sum_i v_i z_i \\ &= \sum_i a_i + \sum_k \max_{\bar{z} \in \mathcal{P}_k} \sum_i v_i z_i. \end{aligned} \quad (17)$$

The last inequality enables us to use the fact that $v_i z_i$ are linear functions, therefore, instead of considering the maximization over the set $\sum_k \mathcal{P}_k$, we can bound $\sum_i v_i z_i$ over each \mathcal{P}_k .

By similar arguments as in the proof of Theorem 4.1, we can prove the following lemma whose proof is provided in [17].

LEMMA 5.2. *For ever k ,*

$$\sum_i U_i'(x_i)x_i^k \geq \frac{1}{1 + 2/\sqrt{3}} \max_{z \in \mathcal{P}_k} \sum_i v_i^k z_i. \quad (18)$$

We use this lemma to prove our main result. On the one hand, if we sum the left-hand side of (18) over all k , we have

$$\sum_{k,i} U_i'(x_i)x_i^k = \sum_i U_i'(x_i)x_i \leq \sum_i U_i(x_i) \quad (19)$$

where the last inequality is true because $U_i(x)$ is a non-negative and concave function for every i . On the other hand, if we sum the right-hand side of (18) over all k , we obtain

$$\frac{\sum_k \max_{z \in \mathcal{P}_k} \sum_i v_i^k z_i}{1 + 2/\sqrt{3}} \geq \frac{\sum_k \max_{z \in \mathcal{P}_k} \sum_i v_i z_i}{1 + 2/\sqrt{3}} \quad (20)$$

where in the last inequality, v_i^k are replaced by v_i , which recall is equal to $\min_k v_i^k$.

Combining (18)–(20), we derive

$$\begin{aligned} \sum_i U_i(x_i) &\geq \frac{1}{1 + 2/\sqrt{3}} \sum_k \max_{z \in \mathcal{P}_k} \sum_i v_i z_i \\ &\Rightarrow (1 + 2/\sqrt{3}) \sum_i U_i(x_i) \geq \sum_k \max_{z \in \mathcal{P}_k} \sum_i v_i z_i. \end{aligned} \quad (21)$$

Finally, from (16), (17) and (21), we have

$$\max_{\bar{z} \in \sum_k \mathcal{P}_k} \sum_i U_i(z_i) \leq (\delta + 1 + 2/\sqrt{3}) \sum_i U_i(x_i)$$

which establishes the asserted result. \square

6. CONVERGENCE AND DISTRIBUTED ALGORITHMS

In this section, we show how a provider may adjust user discrimination weights by an iterative algorithm whose limit points are Nash equilibrium points of the resource competition game studied in earlier sections. We focus on weighted bid auctions but note that similar type of analysis can be carried out for weighted payment auctions. Our aim in this section is to show how such iterative algorithms can be designed in principle.³

³It is beyond the scope of this paper to fully specify implementation details such as online estimation of the elasticity of user utility functions and address the stability in presence of feedback delays.

We replace the polyhedron constraints by incorporating *penalty* function $P(\vec{x})$ in the objective function, which is chosen to confine the allocation vector \vec{x} in the feasible set specified by the polyhedron constraints. Informally, the function $P(\vec{x})$ would assume small values for every feasible allocation vector \vec{x} that is sufficiently away from the boundary of the feasible set and would grow large as the allocation vector \vec{x} approaches the boundary of the feasible set. We assume that $P(\vec{x})$ is a continuously differentiable and convex function. Specifically, we assume that for a collection of functions P_l , one for each of the constraints, we have

$$P(\vec{x}) = \sum_l P_l \left(\sum_j a_{l,j} x_j \right).$$

Indeed, if P_l is a continuously differentiable and convex function for every l , then so is P . We define $V(\vec{x}) = R(\vec{x}) - P(\vec{x})$ where $R(\vec{x})$ is the revenue given by (2). The provider problem is redefined to

$$\text{PROVIDER}' : \text{maximize } V(\vec{x}) \text{ over } \vec{x} \in \mathbb{R}_+^n.$$

We first show how the provider may adjust the user discrimination weights, assuming that the provider knows user utility functions and establish convergence to the Nash equilibrium points in this case. This provides a baseline dynamics that we then approximate as follows. We consider a provider who a priori does not know user utility functions but estimates the needed information in an online fashion while adjusting the discrimination weights. The main idea here is to use an argument based on *separation of timescales* where the provider adjusts the discrimination weights at a slow timescale in comparison with the rate at which bids are received from users, allowing the provider to estimate the needed information about the user utility functions for every given set of discrimination weights. We will formulate iterative algorithms as dynamical systems in continuous time as this is standard in previous work, e.g. [11], and it readily suggests practical distributed algorithms.

User Utility Functions a Priori Known. Suppose that user utility functions are a priori known by the provider (this may be the case if profiles of users are known to the provider, e.g. from the history of previous interactions). The provider announces discrimination weights $\vec{C}(t)$ at every time $t \geq 0$ that are adjusted as follows. The provider computes the allocation vector $\vec{x}(t)$ according to the following system of ordinary differential equations, for some $\alpha > 0$,

$$\frac{d}{dt} x_i(t) = \alpha x_i(t) \frac{\partial}{\partial x_i} V(\vec{x}(t)), \quad i = 1, 2, \dots, n. \quad (22)$$

For every time $t \geq 0$, the provider announces to users the discrimination weights $\vec{C}(t)$ where the discrimination weight for user i is:

$$C_i(t) = x_i(t) + \frac{R(\vec{x}(t))}{U'_i(x_i(t))}.$$

Notice that the right-hand side in (22) requires knowledge of the gradient of the revenue function $R(\vec{x})$, which is given by

$$\frac{\partial}{\partial x_i} R(\vec{x}) = \phi(\vec{x}) \frac{[U'_i(x_i)x_i]'}{(U'_i(x_i)x_i + R(\vec{x}))^2} \quad (23)$$

where

$$\phi(\vec{x}) = \frac{R(\vec{x})}{\sum_j \frac{U'_j(x_j)x_j}{(U'_j(x_j)x_j + R(\vec{x}))^2}}.$$

The convergence to optimal solution of PROVIDER' is showed in the following theorem.

THEOREM 6.1. *Suppose that for every user i , $U'_i(x)$ is a continuously differentiable and concave function and that $P(\vec{x})$ is a strictly convex function. Then, every trajectory $(\vec{x}(t), t \geq 0)$ of the system (22) converges to the maximizer of function $V(\vec{x})$.*

The proof is based on standard application of Lyapunov stability theorem and is thus omitted. It amounts to showing that the function V is a Lyapunov function for system (22), which increases along every trajectory $\vec{x}(t)$, and thus implying convergence to unique maximizer of V .

User Utility Functions a Priori Unknown. We discuss how the user discrimination weights may be adjusted by provider who does not a priori know the user utility functions. The key idea is to use *separation of timescales*: the provider adjusts the discrimination weights at a slower timescale than the timescale at which bids are adjusted by users. Informally speaking, this allows the provider to act as if for every fixed set of discrimination weights, users adjust their bids instantly to the Nash equilibrium bids.

From (6) and the allocation rule $x_i = C_i w_i / \sum_j w_j$, we readily observe that the following identities hold in Nash equilibrium:

$$U'_i(x_i)x_i = \frac{R w_i}{R - w_i} \text{ and } U'_i(x_i)x_i + R = \frac{R^2}{R - w_i}$$

where $R = \sum_j w_j$. Using these identities, we note that the gradient in (23) can be expressed as follows

$$\frac{\partial}{\partial x_i} R = R \frac{\frac{w_i}{R} \left(1 - \frac{w_i}{R}\right) \frac{1}{x_i} + \left(1 - \frac{w_i}{R}\right)^2 \frac{U''_i(x_i)x_i}{R}}{\sum_j \frac{w_j}{R} \left(1 - \frac{w_j}{R}\right)}. \quad (24)$$

Notice that the gradient is fully expressed as a function of the vector of bids \vec{w} except for the term that involves the second derivative of the user utility function. For the case of linear user utility functions, we have $U''_i(x_i) = 0$ for every i , and thus, in this case the gradient of the revenue $R(\vec{x})$ is fully described by the vector of bids \vec{w} .

In general, we assume that at every time $t \geq 0$, the provider sets the user discrimination weight for user i as follows

$$C_i(t) = \frac{R(\vec{w}(t))}{w_i(t)} x_i(t).$$

For the case of linear utility function, $\vec{x}(t)$ is assumed to evolve according to the following system of ordinary differential equations, for $\alpha > 0$ and every $i = 1, 2, \dots, n$,

$$\frac{d}{dt} x_i(t) = \alpha [v_i(w_i(t), R(\vec{w}(t))) - x_i(t) p_i(\vec{x}(t))] \quad (25)$$

where

$$v_i(w_i, R) = R \frac{\frac{w_i}{R} \left(1 - \frac{w_i}{R}\right)}{\sum_j \frac{w_j}{R} \left(1 - \frac{w_j}{R}\right)}$$

$$p_i(\vec{x}) = \sum_l a_{l,i} P'_l \left(\sum_j a_{l,j} x_j(t) \right).$$

Furthermore, we assume natural dynamics for solving USER problem that amounts to adjusting bid $w_i(t)$ by user i according to the following system

$$\frac{d}{dt}w_i(t) = U'_i(x_i(t))x_i(t) - R(\vec{w}(t))\frac{\frac{w_i(t)}{R(\vec{w}(t))}}{1 - \frac{w_i(t)}{R(\vec{w}(t))}}. \quad (26)$$

The convergence for the case of linear user utility functions is established in the following theorem.

THEOREM 6.2. *Suppose that user utility functions are linear. For every sufficiently small $\alpha > 0$, the allocation vector under system (25)-(26) approximates that of the system (22) with an approximation error that diminishes with α .*

The proof is based on applying the *averaging theory* of non-linear dynamical systems [14] and is provided in [17]. It is noteworthy that part of the proof establishes global asymptotic stability of system (26), for the allocation vector $\vec{x}(t)$ fixed to an arbitrary feasible allocation vector \vec{x} for every $t \geq 0$, which is of independent interest as it applies to more general class of utility functions.

Same approach applies more generally to non-linear user utility functions, but the second derivative of the utility function in (24) would need to be estimated in an online fashion from the observed bids submitted by users.⁴ In principle, this can be done by observing the effect of perturbing the allocation for a user on the bid submitted by this user, which we briefly discuss in the following. From (6) and the allocation rule $x_i = C_i w_i / \sum_j w_j$, we have $U'_i(x_i) = \frac{R}{x_i} \frac{w_i}{R} / (1 - \frac{w_i}{R})$. Taking the derivative with respect to x_i , we obtain

$$U''_i(x_i)x_i = \frac{\frac{w_i}{R}}{1 - \frac{w_i}{R}} \left(\frac{\partial}{\partial x_i} R - \frac{1}{x_i} R \right) + R \frac{1}{(1 - \frac{w_i}{R})^2} \frac{\partial}{\partial x_i} \left(\frac{w_i}{R} \right).$$

Plugging this in (24), we have

$$\frac{\partial}{\partial x_i} R = \frac{R}{\sum_j \frac{w_j}{R} (1 - \frac{w_j}{R}) - \frac{w_i}{R} (1 - \frac{w_i}{R})} \frac{\partial}{\partial x_i} \left(\frac{w_i}{R} \right).$$

Therefore, the gradient of the revenue can be fully expressed in terms of the bids \vec{w} and $(\partial/\partial x_i)(w_i/R)$, where the latter term can be estimated in an online fashion by perturbing the allocation of user i and observing the resulting change of w_i/R .

Parking-Lot Example. We demonstrate convergence of the iterative scheme (25)-(26) for the example of parking-lot network that we introduced in Section 3 (Figure 2). Recall that the resource consists of $n \geq 1$ links, each of capacity 1, with a multi-hop user 0 with allocation x_0 at each link and a single-hop user with allocation x_i at link i . User utility functions are assumed to be linear $U_i(x) = v_i x$, for $x \geq 0$, where $v_i = v_1$, for $i = 1, 2, \dots, n$, and $v_0, v_1 > 0$.

The Nash equilibrium allocation and the corresponding revenue are specified by the following lemma whose proof is simple and thus omitted.

LEMMA 6.1. *Let $\eta = \frac{n-1}{2\sqrt{n}} \sqrt{\frac{v_1}{v_0}}$. For the parking-lot scenario, the Nash equilibrium allocation is $1 - x_1$ for user*

⁴Notice that the need to infer the second derivatives of utility functions is intrinsic to the revenue maximization objective and is not an artifact of our auction scheme.




Figure 6: Convergence to equilibrium points for the parking-lot example: (Left) a priori known utility functions and (Right) a priori unknown utility functions.




Figure 7: Another example for the convergence to equilibrium points for the parking-lot: (Left) a priori known utility functions and (Right) a priori unknown utility functions.

0 and x_1 for each user $i = 1, 2, \dots, n$, where

$$x_1 = \begin{cases} \frac{1}{2(1-\eta)}, & \text{for } \eta < 1/2 \\ 1, & \text{for } \eta \geq 1/2 \end{cases}. \quad (27)$$

The revenue in Nash equilibrium is given by

$$R = \begin{cases} \frac{1}{4\eta(1-\eta)}(n-1)v_1, & \text{for } \eta < 1/2 \\ (n-1)v_1, & \text{for } \eta \geq 1/2. \end{cases}$$

We note that in Nash equilibrium, the allocation to any single-hop user is increasing with the ratio of the valuations v_1/v_0 , from $1/2$ for $v_1/v_0 = 0$ to 1 for $v_1/v_0 = n/(n-1)^2$ and beyond this the single-hop user is allocated the entire link capacity.

Numerical examples. We illustrate convergence for two particular cases. In either case, we consider the parking-lot network with $n = 5$ links and linear utility functions specified by $v_0 = 5$ and $v_1 = v_2 = \dots = v_n = 1$. The link cost functions are defined as

$$P'_i(x) = \begin{cases} 0, & 0 \leq x \leq \rho_0, \\ \left(\frac{1}{x} - \frac{1}{1-\rho_0}\right)^p, & \rho_0 < x \leq 1 \end{cases}$$

where $p > 0$ and, in particular, we use $p = 2$ and $\rho_0 = 0.8$. We show results for initial allocation $\vec{x}(0)$ such that $x_1(0) = x_2(0) = \dots = x_n(0)$, so that due to symmetry $x_1(t) = x_2(t) = \dots = x_n(t)$, for every $t \geq 0$. This simplifies the exposition. We validated convergence for various other choices for initial values and other parameters, but for space reasons we confine to the above asserted setting.

In our first case, we consider a closed system with a fixed set of users. In Figure 6, we show trajectories of the allocations $x_0(t)$ and $x_1(t)$ for both a priori known utility functions (left) and a priori unknown utility functions (right),

with α set to 0.1. The results indeed validate convergence to the Nash equilibrium allocation, which are indicated with dashed lines.

Finally, in our second case, we consider an open system where at time $t_1 \geq 0$, a single-hop user departs the system and then another such user arrives at time $t_2 > t_1$. In particular, we use the values $t_1 = 30$, $t_2 = 70$, and $\alpha = 0.8$. Figure 6 well validates convergence to the Nash equilibrium allocation in this case.

7. CONCLUSION

We considered a simple mechanism for allocation of resources that allows for user differentiation and general convex resource constraints. We showed that in a competitive framework where everyone is selfish, the mechanism can guarantee nearly optimal revenue to the provider and competitive social efficiency (including a setting with multiple providers). Besides analysis of equilibrium points of the underlying competition game, we showed how one would design an iterative algorithm that converges to equilibrium points.

The work suggests several interesting directions for future research. First, it would be of interest to study revenue and efficiency properties for classes of user utility functions that are not accommodated in our work. Second, it would be of interest to further explore the space of iterative schemes that are practical and converge to equilibrium points. Finally, one may consider more general settings of multiple competing providers where each user can receive the service only from a subset of providers.

8. REFERENCES

- [1] Amazon Web Services. Amazon EC2 Spot Instances. <http://aws.amazon.com/ec2/spot-instances>, 2010.
- [2] A. Demers, S. Keshav, and S. Shenker. Analysis and simulation of a fair queueing algorithm. *Internetworking: Research and Experience*, 1:3–26, 1990.
- [3] M. Feldman, K. Lai, and L. Zhang. A price-anticipating resource allocation mechanism for distributed shared clusters. In *Proc. of ACM EC*, 2005.
- [4] R. J. Gibbens and F. P. Kelly. On packet marking at priority queues. *IEEE Trans. on Automatic Control*, 47:1016–1020, 2002.
- [5] D. K. Golddenberg, L. Qiuy, H. Xie, Y. R. Yang, and Y. Zhang. Optimizing cost and performance for multihoming. In *Proc. of ACM Sigcomm'04*, Portland, Oregon, USA, 2004.
- [6] A. Gulati, I. Ahmad, and C. A. Waldspurger. Parda: proportional allocation of resources for distributed access. In *Proc. of FAST'09*, San Francisco, CA, 2009.
- [7] B. Hajek and S. Yang. Strategic buyers in a sum bid game for flat networks. In *IMA Workshop*, March 2004.
- [8] J. Hartline and A. Karlin. *Algorithmic Game Theory*, chapter Profit Maximization in Mechanism Design, pages 331–362. editors N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, Cambridge University Press, 2007.
- [9] R. Johari and J. N. Tsitsiklis. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research*, 29(3):402–435, 2004.
- [10] F. Kelly. *Mathematics Unlimited - 2001 and Beyond*, chapter Mathematical modelling of the Internet, pages 685–702. editors B. Engquist and W. Schmid. Springer-Verlag, Berlin, 2001.
- [11] F. P. Kelly. Charging and rate control for elastic traffic. *European Trans. on Telecommunications*, 8:33–37, 1997.
- [12] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan. Rate control for communication networks: Shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49, 1998.
- [13] P. Key, L. Massoulie, and M. Vojnović. Farsighted users harness network time-diversity. In *Proc. of IEEE Infocom 2005*, Miami, FL, USA, 2005.
- [14] H. K. Khalil. *Nonlinear Systems*. Prentice Hall, 3 edition, 2001.
- [15] R. T. Ma, D. M. Chiu, J. C. S. Lui, V. Misra, and D. Rubenstein. On resource management for cloud users: A generalized kelly mechanism approach. Technical report, Technical Report, Electrical Engineering, May 2010.
- [16] T. Nguyen and E. Tardos. Approximately maximizing efficiency and revenue in polyhedral environments. In *Proc. of ACM EC'07*, pages 11–19, 2007.
- [17] T. Nguyen and M. Vojnović. The weighted proportional allocation mechanism. Technical Report MSR-TR-2010-145, Microsoft Research, 2010.
- [18] A. Ozdaglar and R. Srikant. *Algorithmic Game Theory*, chapter Incentives and Pricing in Communication Networks. editors N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, Cambridge University Press, 2007.
- [19] A. K. Parekh and R. G. Gallager. A generalized processor sharing approach to flow control in integrated services networks: The multi node case. *IEEE/ACM Trans. on Networking*, 2:137–150, April 1994.
- [20] S. Shakkottai and R. Srikant. Economics of network pricing with multiple isps. *IEEE/ACM Trans. on Networking*, 14(6):1233–1245, 2006.
- [21] S. Shakkottai, R. Srikant, A. Ozdaglar, and D. Acemoglu. The price of simplicity. *IEEE Journal on Selected Areas in Communications*, 26(7):1269–1276, 2008.
- [22] J. Tirole. *The Theory of Industrial Organization*. The MIT Press, 2001.
- [23] G. Wang, A. Kowli, M. Negrete-Pincetic, E. Shafieepoorfard, and S. Meyn. A control theorist's perspective on dynamic competitive equilibria in electricity markets. In *Proc. 18th World Congress of the International Federation of Automatic Control (IFAC)*, Milano, Italy, 2011.
- [24] Wikipedia. California electricity crises. http://en.wikipedia.org/wiki/California_electricity_crisis, 2010.